

AD-A041 907

MARYLAND UNIV COLLEGE PARK COMPUTER SCIENCE CENTER
A DISCUSSION OF HARDWARE IMPLEMENTATION AND FABRICATION FOR AN --ETC(U)
JAN 77 T WILLETT, N BLUZER

F/G 9/3

DAAG53-76-C-0138

NL

UNCLASSIFIED

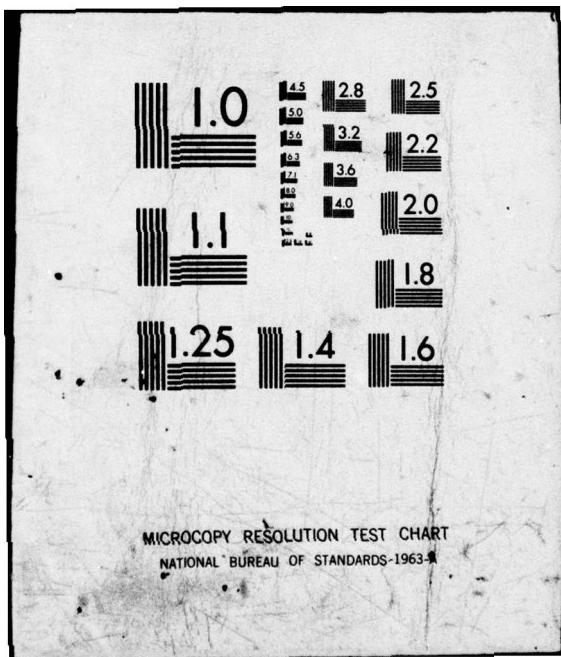
OF
ADAD041 907



END

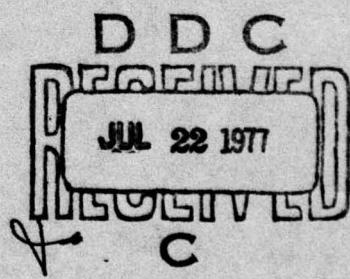
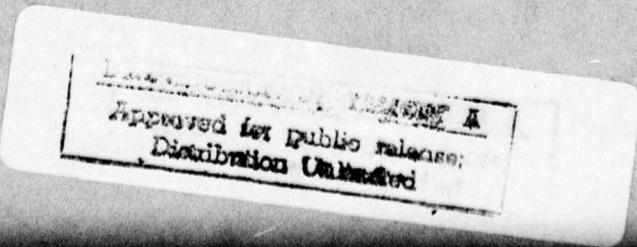
DATE
FILMED

8 - 77



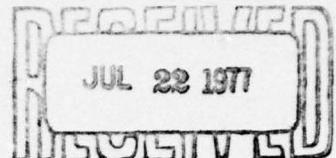
AD No. _____
DDC FILE COPY

ADA041907



B-51
①

D D C



(6) A DISCUSSION OF HARDWARE
IMPLEMENTATION AND FABRICATION
FOR AN AUTOMATIC TARGET CUEING SYSTEM.

January 31, 1977

Q Quarterly status report no 3

(13) This is the third quarterly status report
on a program for Recognition Technology for a
Smart Sensor, conducted by Westinghouse for
Maryland under Contract DAAG 53-76-C-0138 with
the U. S. Army Electronics Command, Night Vision
Laboratory, Ft. Belvoir, Va. 22060.

(10) Thomas J. Willett
Nathan Blazer

(11) 31 Jan 77

Prepared for

Computer Science Center
University of Maryland
College Park, Maryland 20742

(12) 33 p.

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

By

Westinghouse Defense and Electronic Systems Center
Systems Development Division
Baltimore, Maryland 21203

103-018

7/B

TABLE OF CONTENTS

INTRODUCTION

1.0 SYSTEM FLOW

1.1 Algorithms

- 1.1.1 Median Filter
- 1.1.2 Gradient Operator
- 1.1.3 Non Maximum Suppression
- 1.1.4 Gray Level Threshold Determination
- 1.1.5 Connected Components
- 1.1.6 Perimeter Match

1.2 Data Flow

1.3 Storage Requirements

2.0 HARDWARE IMPLEMENTATION

2.1 Median Filter

- 2.1.1 CCD Charge Quantizer
- 2.1.2 CCD Sorter

2.2 Non Maximum Suppression

3.0 HARDWARE FABRICATION

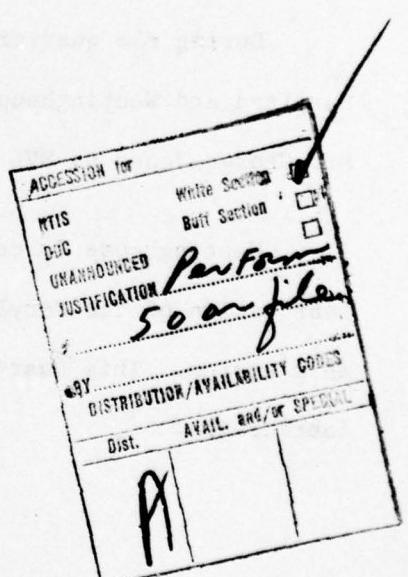
3.1 Gradient Operator

3.2 Median Filter

3.3 Serpentine Delay

3.4 General Observations

4.0 FOCAL PLANE AREA



INTRODUCTION

This is the third quarterly status report on a program for RECOGNITION TECHNOLOGY FOR A SMART SENSOR, being conducted by Westinghouse Systems Development Division for the Computer Science Center, University of Maryland. The program consists of three phases, as follows:

Phase I Task and Technology Review (3 months)

Phase II Algorithm Selection and Test (9 months)

Phase III Hardware Development (9 months)

This report covers the second 3 months of the Phase II effort. The report was prepared by Mr. Thomas Willett and Dr. Nathan Bluzer of Westinghouse. The Westinghouse program manager is Dr. Glenn E. Tisdale.

During the quarter, five meetings were held between members of the Maryland and Westinghouse teams. Mr. John Dehne, NVL program manager, and Mr. George Jones of NVL attended several of the meetings.

Westinghouse is concentrating on the hardware implementation and fabrication of the Maryland algorithms for the focal plane and treating them as a system. This quarter marks the shift in emphasis from implementation to fabrication.

1.0 SYSTEM FLOW

This section describes a preferred set of algorithms developed by Maryland which tentatively comprises the first portion of a cueing system. A system flow chart is shown in Figure 1-1. A description of data flow and storage requirements is included.

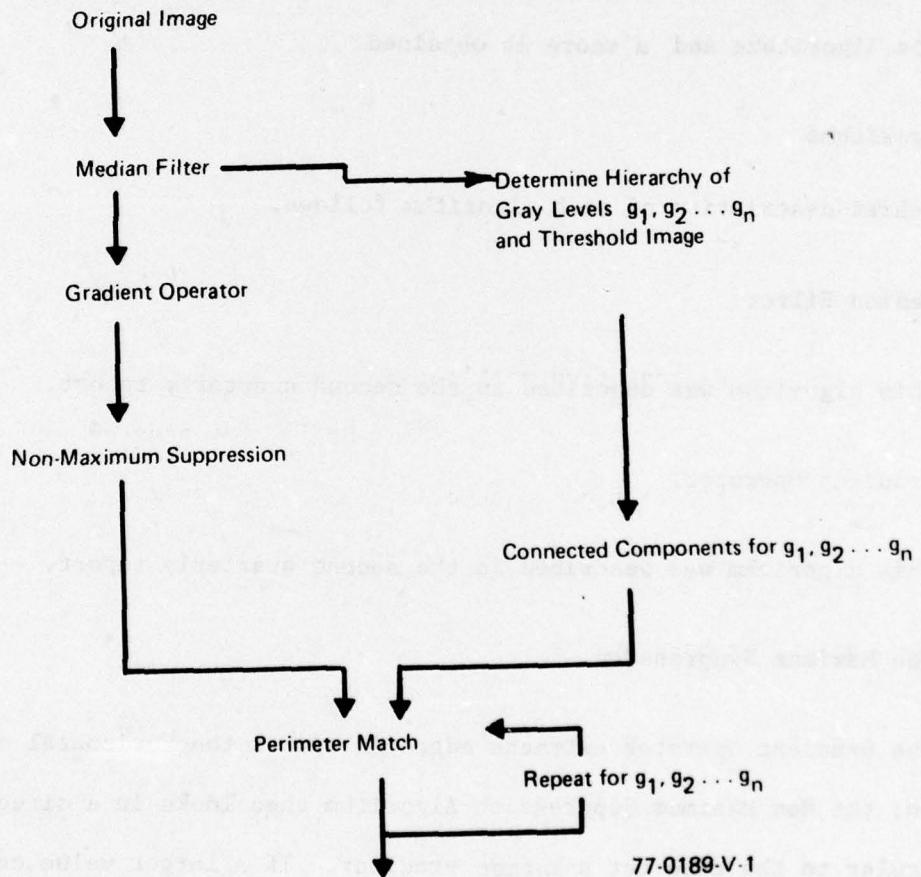


Figure 1-1. System Flow Chart

In general, the Median Filter acts to suppress noise. The Gradient Operator extracts edges; the width of these edges is reduced by the Non Maximum Suppression Algorithm. At the same time a set of gray levels (g_1, g_2, \dots, g_n) is determined from the entire frame. The filtered image is thresholded at each gray level and a Connected Components Algorithm partitions the thresholded image into regions. A Matching Algorithm correlates perimeter points formed independently by the Non Maximum Suppression and Connected Components Algorithms and a score is obtained.

1.1 Algorithms

A short description of each algorithm follows.

1.1.1 Median Filter

This algorithm was described in the second quarterly report.

1.1.2 Gradient Operator

This algorithm was described in the second quarterly report.

1.1.3 Non Maximum Suppression

The Gradient Operator extracts edges in either the horizontal or vertical direction; the Non Maximum Suppression Algorithm then looks in a direction perpendicular to the edge for a larger gradient. If a larger value cannot be found, the edge under consideration is retained; the edge is removed if a larger value is found. The Algorithm is shown in Figure 1-2.

X	X	X		
X	X	X	If any X > Y, Y = 0	
	X		otherwise retain Y	
Y		→	X, Y are gradient values	
	X			
X	X	X		
X	X	X		

77-0189-V-2

Figure 1-2. Non-Maximum Suppression

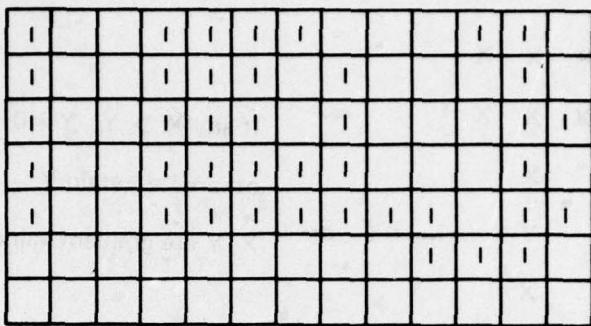
The gradient under consideration is a horizontal one and the area examined for larger gradients is in the vertical direction. This imposes an additional requirement on the Gradient Operator in that the direction of the larger gradient must be retained.

1.1.4 Gray Level Threshold Determination

Maryland presently is evaluating several approaches to this determination; hardware implementation will be held in abeyance until completion.

1.1.5 Connected Components

Assume that the original image has been thresholded and the result is in binary form with gray levels exceeding g_1 shown as 1's in Figure 1-3.



77-0189-V-3

Figure 1-3a. Binary Image

A		B	B	B	B		D	D
A		B	B	B	C		D	
1		1	1	1	1		1	1
1		1	1	1	1		1	
1		1	1	1	1	1	1	1
						1	1	1

77-0189-V-4

Figure 1-3b. Computations for Second Row

Two image lines are retained in memory so that each pixel can examine its neighbors to the left and right and above and below. No diagonal connections are permitted under this convention, and an adjacent (horizontal or vertical) pixel must be occupied in order to make a connection. No skips or gaps are allowed, and the computations start one pixel in from the edge. In Figure 1-3b, there are four distinct regions, A, B, C, and D. The only possible connection between regions B and C is through a diagonal, which is not allowed.

Computations for the fourth row are seen in Figure 1-3c.

A		B	B	B	B			D	D	
A		B	B	B		C			D	
A		B	B	B		C			D	D
A		B	B	B	C	C			D	
I		I	I	I	I	I	I	I	I	I
							I	I	I	

77-0189-V-5

Figure 1-3c. Computations for Fourth Row

Here, there is a connection between regions B and C and an equivalence statement, $B = C$, is carried along. At the end of the sixth row, there is another connection between C and D ($C = D$) and all the regions are completed as seen in Figure 1-3d.

A		B	B	B	B			B	B	
A		B	B	B		B			B	B
A		B	B	B		B			B	
A		B	B	B		B			B	
A		B	B	B	B	B	B	B	B	B
							B	B	B	

77-0189-V-6

Figure 1-3d. Completed Image

The areas of A, B, C and D are computed by cumulating the number of pixels assigned to each. The perimeter is calculated by cumulating the number of pixels assigned to each region which are neighbors of zeros, i.e., the neighbors did not exceed the gray level threshold, g_1 .

1.1.6 Perimeter Match

Suppose an image has been processed through the Non Maximum Suppression and Connected Components Algorithms independently as seen in Figure 1-4.

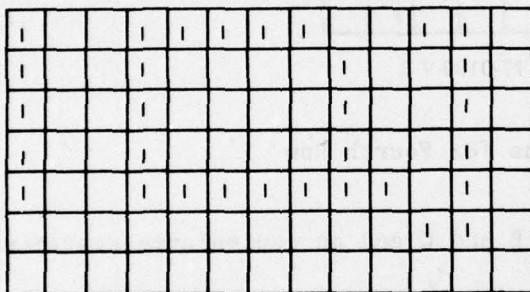


Figure 1-4a. Non Maximum Suppression Output

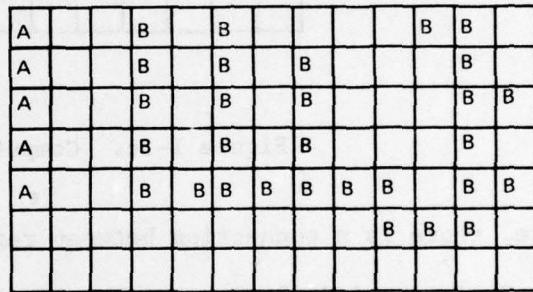


Figure 1-4b. Connected Component Output

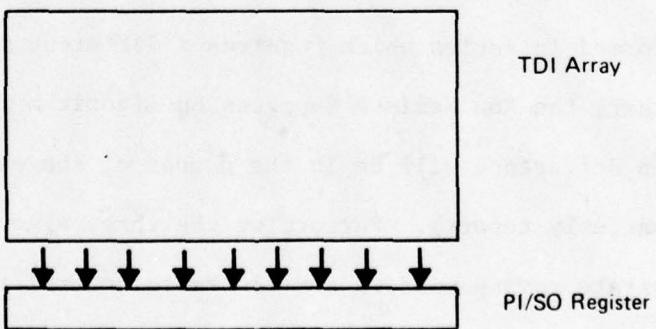
Region A has a perfect score: every pixel of maximum gradient is matched as a perimeter point of A. Region B matched 23 points out of 26 possible gradient points but it also produced 5 perimeter points which were not matched by gradient pixels.

1.2 Data Flow

The Median Filter and Gradient Operator are calculated for a small window which moves over the entire frame. The technique for extracting the appropriate pixels from the focal plane for each window position was described in the second

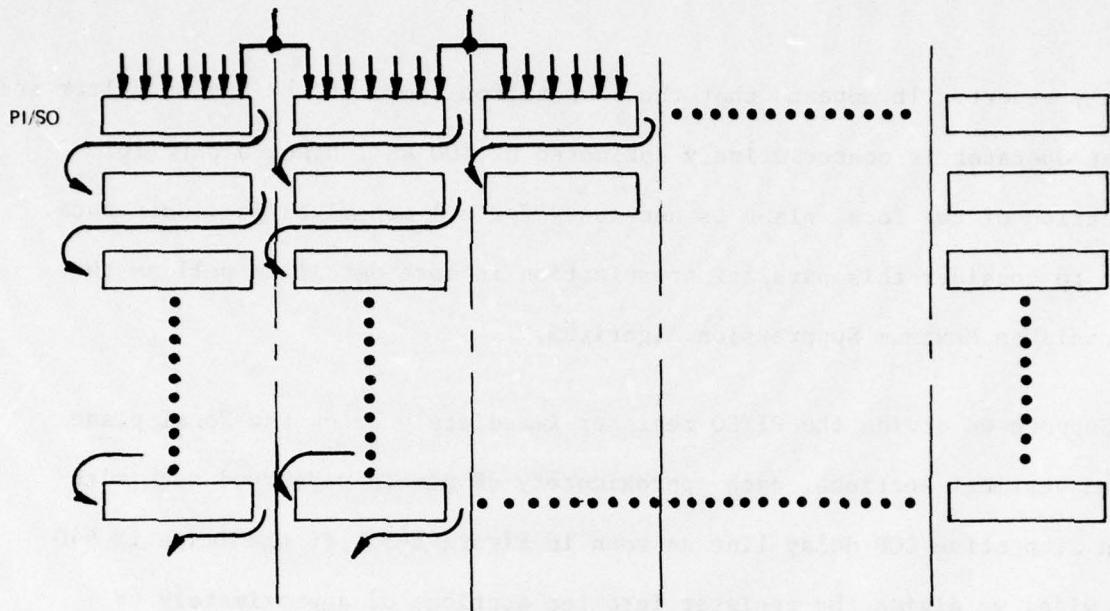
quarterly report. It appears that the computation speed of the Median Filter and Gradient Operator is conservatively estimated at 100 kHz, hence a parallel organization of the focal plane is necessary for a 1 megapixel/sec. data rate. We want to consider this parallel organization in more detail as well as the additional Non Maximum Suppression Algorithm.

Suppose we divide the PI/SO register immediately below the focal plane into ten vertical sections, each approximately 68 pixels wide, and each with its own serpentine CCD delay line as seen in Figure 1-5. If the image is 640 pixels wide, we divide the register into ten sections of approximately 68 pixels each to avoid problems associated in calculating medians and gradients along the edge of an image. Each vertical section is eight stages long to accommodate the Gradient Operator, which requires eight lines of storage, and 68 pixels wide for a total number of 544 shifts at 100 kHz. This appears to avoid numerical integrity problems.



77-0189-V-8

Figure 1-5a. Focal Plane Arrangement



77-0189-V-9

Figure 1~5b. PI/SO and Serpentine in Parallel

However, the system flow chart in Figure 1-1 differs from the second quarterly report in regard to the order of performance of the algorithms. There the Median Filter and Gradient Operator were performed in parallel, and here they are performed in series which requires a different moving window arrangement. Further, the Non Maximum Suppression Algorithm is next in the sequence. The main difference will be in the number of stages of the serpentine (see the second quarterly report). Performing the three algorithms consecutively requires three separate moving windows and clocks to control the non destructive readouts for each. The Median Filter Algorithm requires 5 stages of delay, the Gradient Operator requires 8 lines (stages), and the Non Maximum Suppression

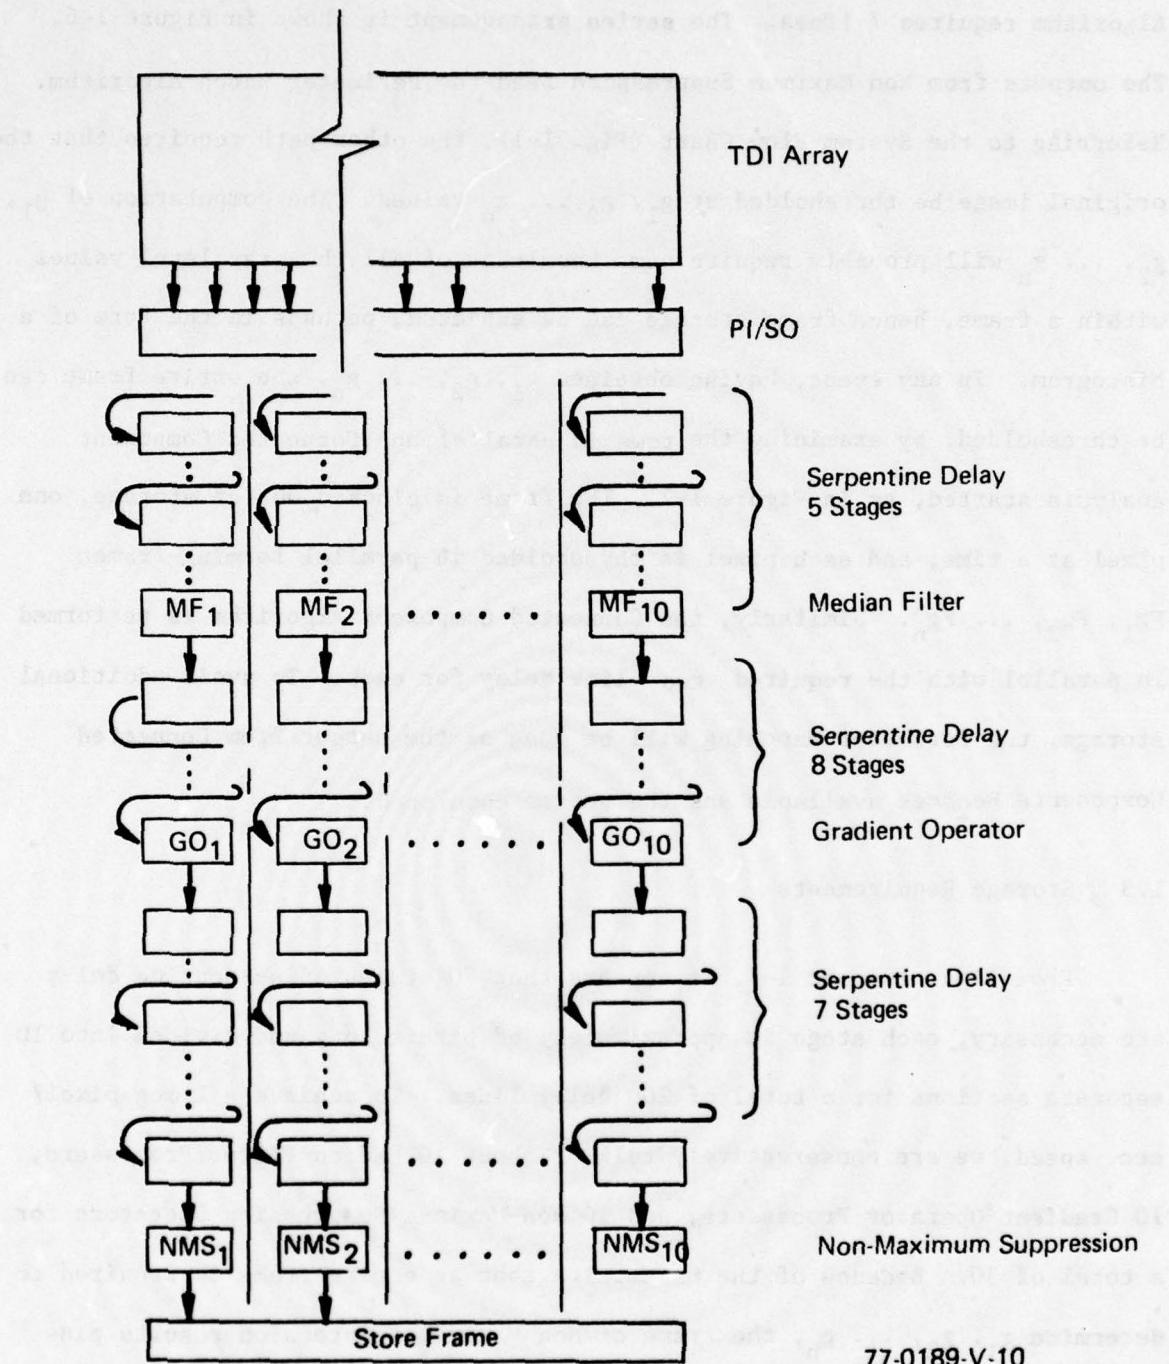
Algorithm requires 7 lines. The series arrangement is shown in Figure 1-6.

The outputs from Non Maximum Suppression feed the Perimeter Match Algorithm.

Referring to the System Flow Chart (Fig. 1-1), the other path requires that the original image be thresholded at g_1 , $g_2 \dots g_n$ values. The computation of g_1 , $g_2, \dots g_n$ will probably require some knowledge of all the gray level values within a frame, hence frame storage can be expected, perhaps in the form of a histogram. In any event, having obtained g_1 , $g_2, \dots g_n$, the entire frame can be thresholded, by examining the rows in parallel, and Connected Component analysis started, as in Figure 1-7. The frame is clocked out of storage, one pixel at a time, and each pixel is thresholded in parallel forming frames Fg_1 , $Fg_2, \dots Fg_n$. Similarly, the Connected Component Algorithm is performed in parallel with the required two line delay for each. To avoid additional storage, the Perimeter Matching will be done as the output from Connected Components becomes available and the scores cumulated.

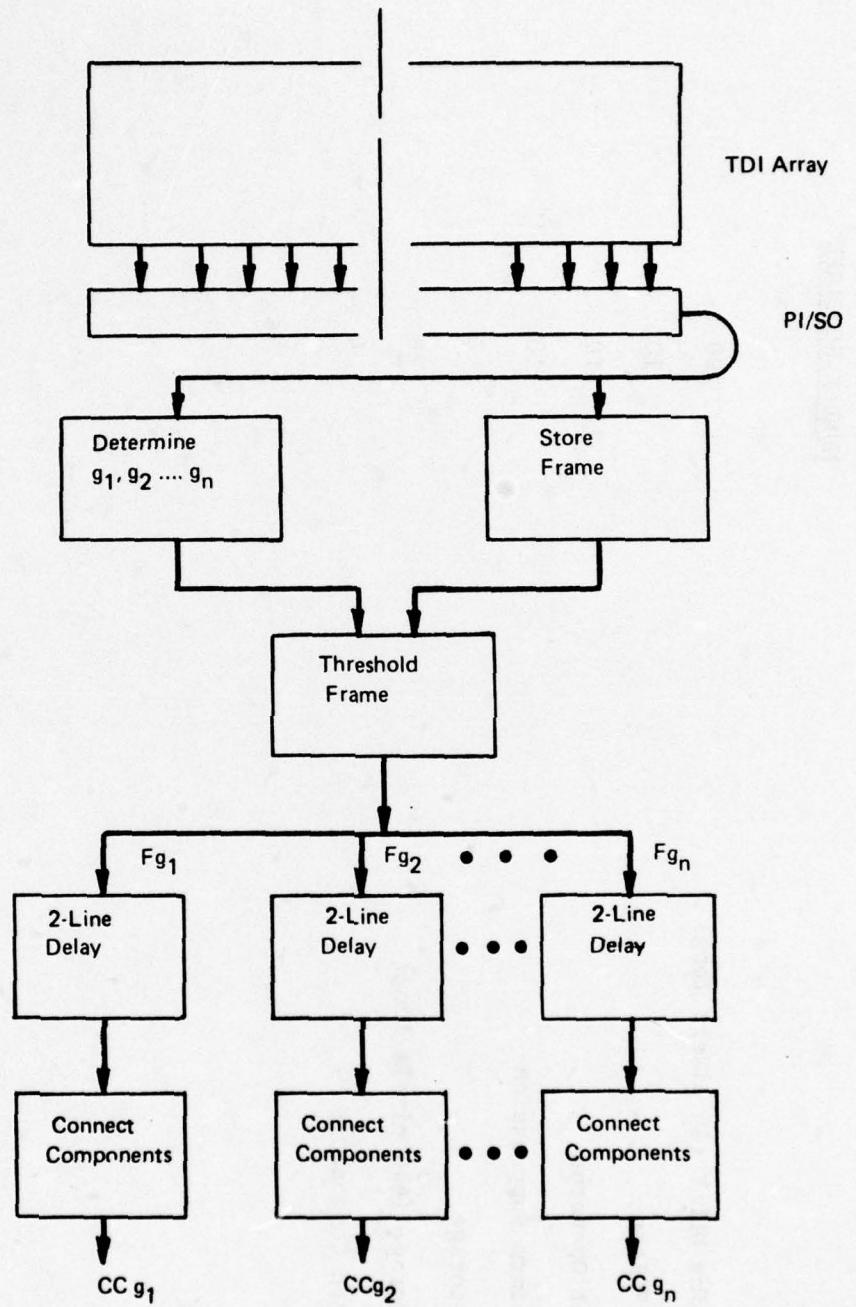
1.3 Storage Requirements

From Figs. 1-6 and 1-7, it appears that 20 stages of serpentine delay are necessary, each stage is approximately 68 pixels long and divided into 10 separate sections for a total of 200 delay lines. To achieve a 1 megapixel/sec. speed, we are conservatively talking about 10 Median Filter Processors, 10 Gradient Operator Processors, and 10 Non-Maximum Suppression Operators for a total of 30. Because of the assumption that an entire frame is required to determine g_1 , $g_2, \dots g_n$, the frame of Non Maximum Suppression results plus the original frame of gray levels must be stored. Assuming four thresholds, as an example, an additional 12 lines of delay are required, plus four Connected Component Processors. It should be noted that these are tentative results with final conclusions reserved for the next quarter, particularly with regard to Connected Components. The results are summarized in Table 1-1.



77-0189-V-10

Figure 1-6. Focal Plane Data Flow



77-0189-V-11

Figure 1-7. Connected Components Data Flow

TABLE 1-1. NUMBER OF PROCESSORS REQUIRED FOR 1 MEGAPIXEL/SEC.

<u>PROCESS</u>	<u>NUMBER REQUIRED</u>
Serpentine Delay (68 pixels long)	200
Median Filter	10
Gradient Operator	10
Non Maximum Suppression	10
Frame Storage	2
Line Storage (640 pixels long)	8
Connected Components	4

2.0 HARDWARE IMPLEMENTATION

In the prior section, we discussed system flow, algorithms, data flow, and storage requirements for the Maryland design. In this section, we shall discuss specific hardware techniques to perform the algorithms.

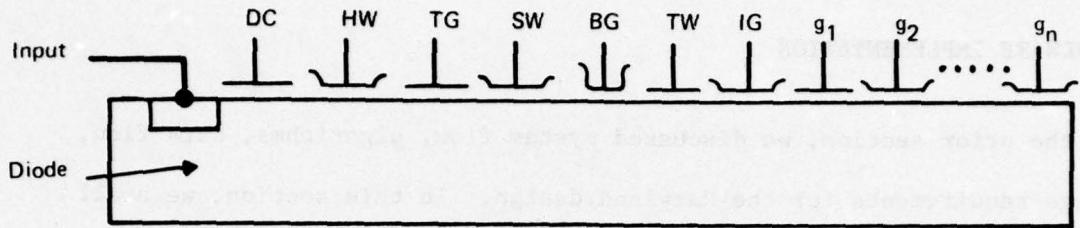
In the second quarterly report, an implementation of the Median Filter was described; that work is extended and completed here. The sorter analyzed in the Median Filter section is applicable to the Non Maximum Suppression Algorithm which is discussed in the next section.

2.1 Median Filter

2.1.1 CCD Charge Quantizer

In signal processing applications the need arises for quantizing the signal (e.g., discrete representation) to facilitate special processing. Such special processing operations include Median Filtering, A/D conversion, sorting and comparisons on the focal plane. A novel technique for signal quantization under the constraints of low power, small space, and cryogenic temperatures is to employ a special charge coupled device. With a proper structure and operating clocks, a CCD can be employed to transform an analog signal charge packet into a thermometer representation (see second quarterly report). This representation can be employed as an intermediate step in converting analog signals to a binary code.

Consider a CCD charge quantizer shown in Figure 2-1. The device is structured to transform the signal, S , into a charge packet, Q , and is insensitive to d.c. threshold effects. A charge signal, C , is injected into



77-0189-V-12

Figure 2-1. Charge Quantizer

the holding well designated by HW. The diode is used as a charge source and the DC gate acts as a blocking electrode and a scuppering electrode for HW.

Eliminating DC offsets requires a calibration of the holding well depth. The charge C in the holding well is scuppered by lowering the Transfer Gate (TG) potential while the Signal Well (SW) is enabled attractive. The charge scuppered (removed) from the C charge is also removed from the CCD register.

Having calibrated the holding well, the voltage on the Transfer Gate (TG) is changed by the signal voltage, S, such that the potential barrier formed by the Transfer Gate is lowered causing charge to spill into the Signal Well, SW. The signal in the SW location is equal to Q, which is proportional to the signal S, and is independent of threshold variations.

The signal charge Q located at SW is removed by repeatedly taking out a charge quantum q. The value of q is determined by the voltage between the Blocking Gate (BG) and the Thimble Well (TW) and is adjusted to reflect the number of grey resolution levels desired.

Each measured charge quantum q removed from the Signal Well is shifted into the CCD shift register by gates G₁ through G_n. The number of locations N

in the CCD shift register is equal to the number of gray level resolution elements required. The quanta are shifted into a serial in/parallel out (SI/PO) CCD shift register. After N shift periods, only K ($K \leq N$) locations in the CCD shift register will be filled, each with a quantum q . The remaining $N-K$ locations will be empty of charge as in Figure 2-2. The register contents are next shifted in parallel out of the serial CCD into a memory module. Repeated operation of the CCD quantizer will result in converting an analog data string into a column array where the size of each column represents the amplitude of a single analog data element in the data string.

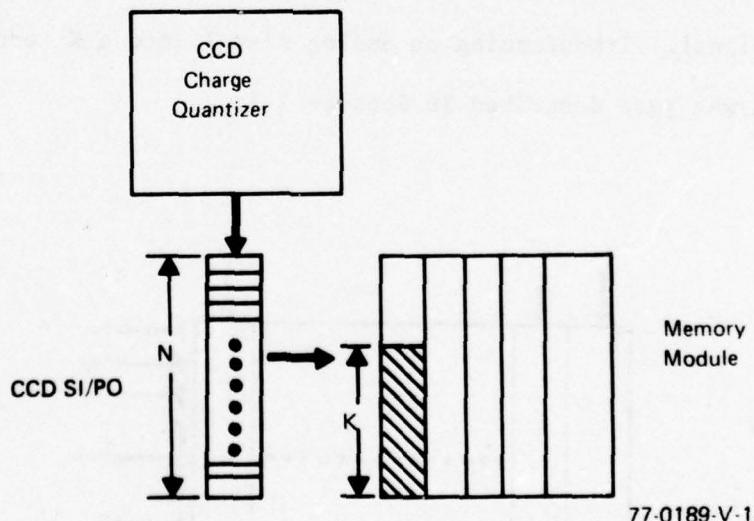


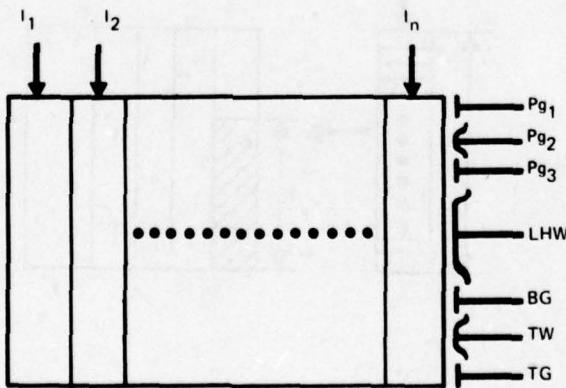
Figure 2-2. CCD Charge Quantizer

2.1.2 CCD Sorter

Although this sorter is included in the Median Filter section, it is directly applicable to the Non Maximum Suppression Algorithm.

Many applications require sorting N data elements according to size, the largest element first and the smallest element last. We will describe a CCD sorter which requires only M sorting operations to rearrange M data elements located in random order. Conventionally, $\frac{M(M-1)}{2}$ comparisons are required for sorting M elements.

The CCD module we propose is shown in Figure 2-3, where the N parallel input channels are equal to the maximum number of quantized grey levels comprising the signal. Transforming an analog signal into a K resolution thermometer code was just described in Section 2-1.



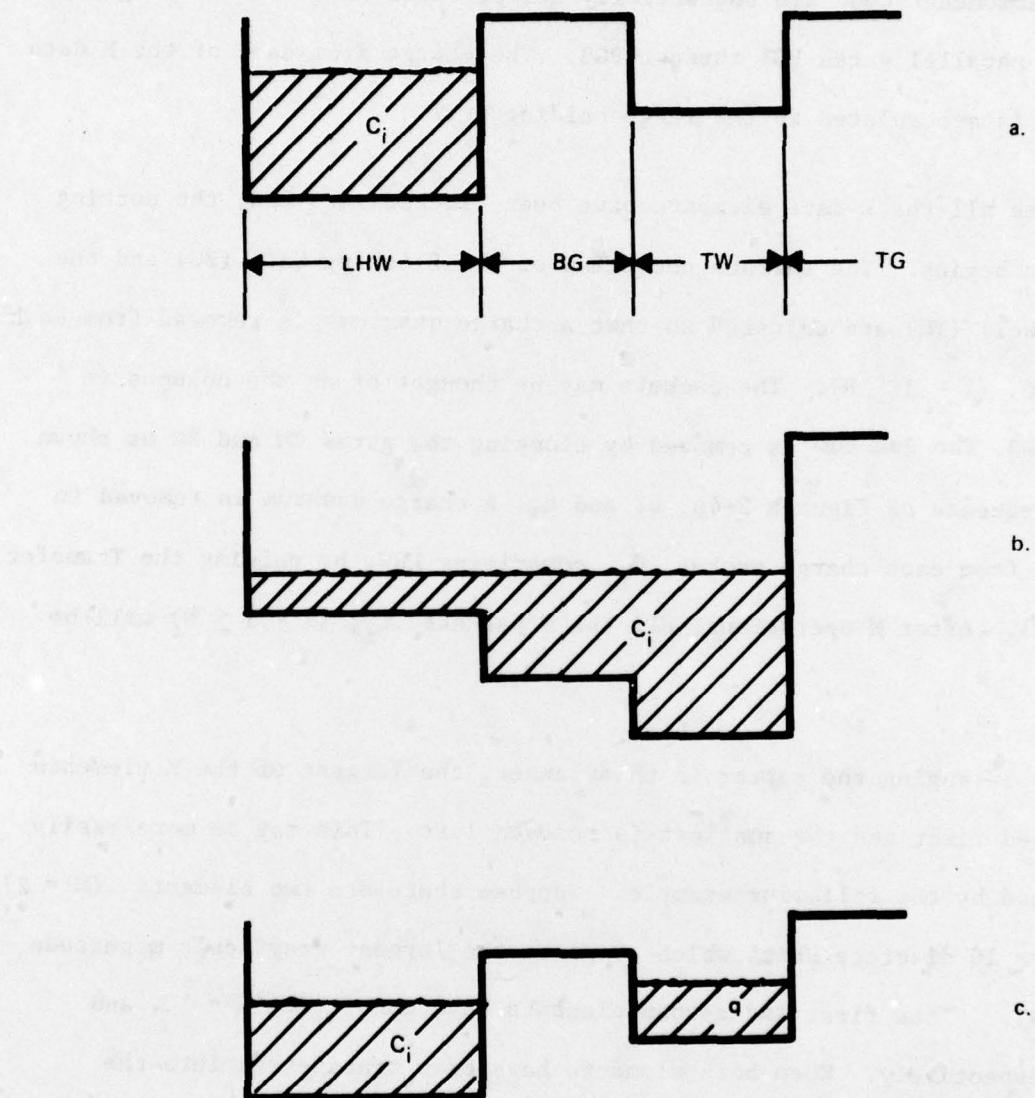
770189-V-14

Figure 2-3. CCD Sorter

To perform the sorting, all the M data elements represented by an N level thermometer code are sequentially shifted into the large holding well (LHW) by parallel gates PG1 through PG3. The charge from each of the M data elements is accumulated in the large holding well.

When all the M data elements have been clocked into LHW, the sorting operation begins. The surface potential of the Blocking Gate (BG) and the Thimble Well (TW) are adjusted so that a charge quantum q is removed from each packet, C_i ($1 \leq i \leq N$). The packets may be thought of as the columns in Figure 2-3. The quantum is removed by clocking the gates TW and BG as shown in the sequence of Figures 2-4a, b, and c. A charge quantum is removed in parallel from each charge packet, C_i , comprising LHW, by pulsing the Transfer Gate (TG). After M operations, all the N packets, C_i , ($i \leq 1 \leq N$) will be empty.

By arranging the sorter in this manner, the largest of the M elements is removed first and the smallest is removed last. This may be more easily understood by the following example. Suppose there are two elements ($M = 2$), there are 10 discrete units which comprise the largest grey scale magnitude ($N = 10$), the first and second elements have magnitudes $N = 10$, and $N = 1$ respectively. When both elements have been transferred into the LHW, the contents of the charge packets, C_i , are $C_1 = q$, $i = 1, 2, \dots, 9$ and $C_{10} = 2 q$. To remove the quanta, all C_i are pulsed and an amount q is removed from all the packets. This is equivalent to removing the largest element first. The contents of the sorter are now $C_1 = 0$, $i = 1, 2, \dots, 9$ and $C_{10} = q$. Pulsing the sorter again permits removal of the smaller element. The results can be stated more generally: the elements clocked out of the CCD



77-0189-V-15

Figure 2-4 a,b,c. Charge Removal Sequence

sorter will decrease monotonically in size with the exception of equal values.

2.2 Non Maximum Suppression

One approach to obtaining the outline of an object is by extracting only the largest gradients in a scene. To obtain the largest gradients implies that each gradient must be compared with those in its neighborhood; see Section 1.3 for a description of the Non Maximum Suppression Algorithm.

Embodiment of the Non Maximum Suppression Algorithm (NMS) requires several operations with CCD structures; the types of operations can be determined by decomposing the NMS. A key part of NMS is extracting the largest X_m gradient value in the neighborhood surrounding y ; X_m is then compared to the gradient value y_g representing the y th pixel. Sorting the X_i values to obtain X_m can be accomplished by the sorting operator described earlier. Hence the only operation required for implementing NMS is a method for comparing X_m to y_g according to the rule described in Section 1.1.3. Consider the block diagram of Figure 2-5.

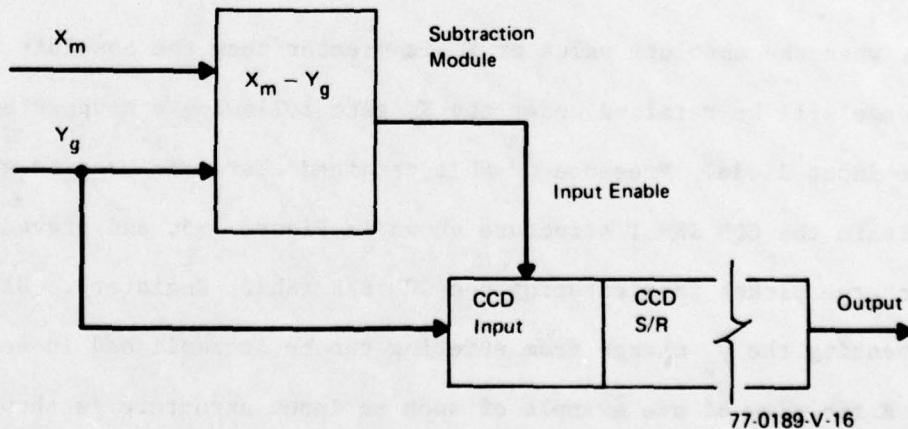


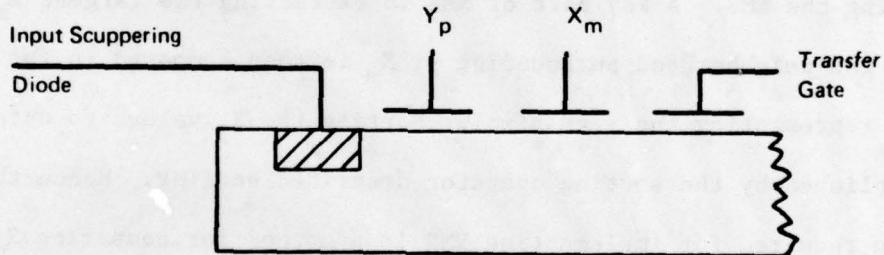
Figure 2-5. Partial Block Diagram of NMS

The subtraction module has two inputs, one from the sorter (X_m) and the other from the gradient value y_g in the y th pixel location. The Subtraction Module is a CCD structure into which X_m and y_g are connected such that the output (INPUT ENABLE SIGNAL) will be

$$0 \text{ if } y_p \leq X_m$$

$$1 \text{ if } y_p > X_m.$$

Such a function can be accomplished if we connect the y_p and X_m inputs as shown in Figure 2-6 and employ scuppering type CCD injection.



77-0189-V-17

Figure 2-6. Input Structure of the Subtraction Module

Clearly when the absolute value of X_m is greater than the absolute value of y_p , charge will be retained under the X_m gate following a scuppering operation via the input diode. Presence of this retained charge is used to enable a gate within the CCD INPUT structure shown in Figure 2-5, and prevent the y_p input charge packet from entering the CCD S/R (Shift Register). Blocking or preventing the y_p charge from entering can be accomplished in several ways. A top view of one example of such an input structure is shown in Figure 2-6.

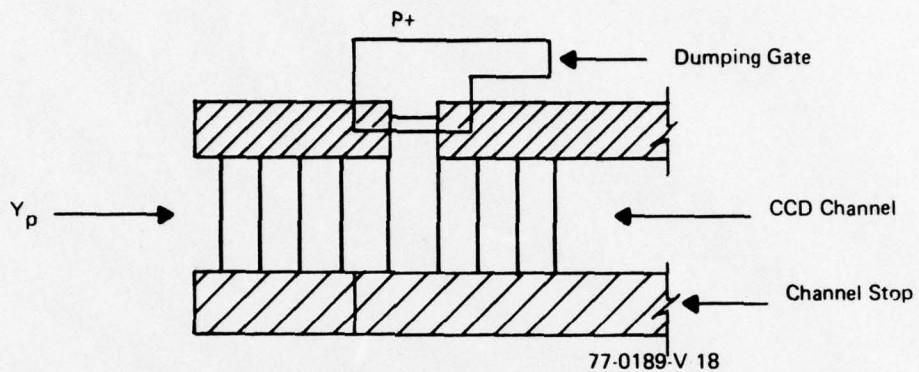


Figure 2-6. Input Blocking Structure

The injected charge y_p is clocked from left to right in the CCD channel. When it reaches the third CCD gate position, a dumping gate is enabled (blocked) if the output from the subtraction module is one (zero). If a "one" is present from the subtraction module, the dumping gate opens a channel for the charge located beneath the third gate to free a P + (N +) charge sink if we employ a surface P (N) channel CCD.

3.0 HARDWARE FABRICATION

In previous work, we have described the hardware implementation of a number of algorithms. In this section we shall take that work a step further and consider the fabrication of these implementations. Specifically, we shall consider chip size, cryogenic problems, speeds, yields, and power consumption relevant to the Gradient Operator, the Median Filter, and the Serpentine Memory.

3.1 Gradient Operator

A major assumption in the analysis is that the inputs for the Gradient Operator (i.e., A, B, C, and D) will be obtained from a separate IC chip which will be a part of the serpentine memory module. The structure of this module will be addressed in another section.

The size of the Gradient Operator chip will be deduced by assigning real estate to each operation performed by the Operator. A key operation is the absolute subtraction module (ASM) which obtains the absolute difference between two inputs and yields a charge representing that quantity. Each difference CCD structure will nominally require a channel 1.2 mils wide; four input channels are needed to provide four charge packets, two representing $|A_i - B_i|$ and two representing $|C_i - D_i|$. The length of each ASM will be 4 mils, a size sufficient to provide a readout structure necessary to drive the second stage of the Operator. The second stage selects which output $|A_i - B_i|$ or $|C_i - D_i|$ is the largest gradient of the i th pixel location. Combining the real estate requirement for the first and second stages, we calculate a chip size of 8 mils x 10 mils.

We assume a four phase gate construction; a smaller number of phases (which requires less chip area) could be used; however, speed-charge handling capacity and ease of fabrication favors four phase construction.

The structure advocated is exclusively based on MOS FET and CCD technology. Both MOS FET and CCD structures exhibit improved performance at cryogenic temperatures greater than 30°K. At very low cryogenic temperatures ($\geq 30^{\circ}\text{K}$), the performance of MOS CCD structures begins to show significant degradation. Relative to room temperature performance, experiments have shown that with cryogenic temperatures we should obtain higher operational speeds and lower noise figures. The improved performance is attributed to increases in mobility resulting from lower levels of phonon scattering of the signal carriers.

The fabrication yield depends on chip size and the number of steps. The process is very similar to that for making surface channel CCD and we estimate six photolithographic masks. Cognizant of these similarities, we expect a yield of better than 50%. The variables which will influence the final chip configuration will be speed, charge handling requirements and resolution (uniformity). Our present design is conservatively aimed at a speed of 100 kHz. Higher speeds are possible, but increasing the operating speed from 100 kHz to 1 MHz will require special and more difficult structures.

Power consumption of the Gradient Operator will depend on the operating frequency and checking voltages. Conventionally, the power consumed by a CCD type structure is expressed as

$$P = CV^2 f N$$

where N is the number of gates, C is the capacitance of each gate, V is the clocking voltage and f is the operating frequency. Computing the power requirements we obtain less than 10 milliwatts. This level of power consumption is exclusive of the power requirements of the clocking circuitry required to operate the Gradient Operator.

3.2 Median Filter

In the second quarterly report, Maryland reported on the significance of the Median Filter Operator and Westinghouse described an embodiment. In this section, we shall consider aspects pertinent to fabrication.

The MFO chip as considered below will not include peripheral clocking circuits or a structure for summing the output from the serpentine CCD delay. We assume an MFO operating on 25 pixels located within a moving window; provisions for obtaining the 25 pixels will be built into the CCD serpentine delay structure in the form of non-destructive readouts. Each data element (pixel) is assumed to have a dynamic range equivalent to a 32 level grey scale.

The size of the chip is determined primarily by the number of pixels and grey levels. The proposed MFO is required to operate as a moving window device which requires a CCD memory capable of storing and shifting 25 data elements each of which is quantized within a 32 level grey scale. A bank of CCD memory registers with 25×32 storage locations can be achieved by a 64 mil by 64 mil module. Included in this estimate are areas for incorporating output and input structures to the CCD memory.

Another major block of the MFO is the sorting module in which the data elements are arranged according to size. This requires a bank of 32 CCD shift registers which are 25 elements long and each row is capable of being independently shifted left or right. An area 100 mils wide by 64 mils long is sufficient.

Finally, the area required for controlling the clocks operating the sorting module is estimated to be 100 mils by 2 mils.

Summing the different component areas comprising the MFO, we arrive at an area estimate of 100 mils by 128 mils.

All the elements used in modelling the MFO are based on field effect phenomena, hence we expect improved performance at cryogenic temperatures in accordance with experimental observations. As described in the Gradient Operator section, we expect improvement in bandwidth and noise reduction.

The size of the MFO (128 x 100 mils) represents a large scale integration device and significant complexity will be encountered during fabrication and test. We calculate that eight mask levels will be required. The yield, dependent on chip size and the number of masking steps, is estimated to be 5%.

Required power will be larger than that consumed by a conventional CCD device; the demand for more power comes from the active logic devices used for clock control (shift left or right) functions. Generally, power consumed depends on operating speed, component cost, and system layout. Assuming an operating speed of 100 kHz, the MFO will require less than 100 milliwatts of power.

3.3 Serpentine Delay

A large number of delay elements are required for focal plane processing; the elements must have large memory capacity, good transfer efficiency, and non destructive readout structures.

The serpentine deployment shown in Figure 1-6 requires a large number of transfers which causes degradation in the modulation transfer function (MTF). This negative effect can be reduced by segmenting the image into several columns, each of which will be processed in parallel with the other columns. Such a segmentation not only reduces the number of CCD transfers per delay element but reduces the operating speed. We postulated an algorithm operating speed of 100 kHz based on hardware implementations. This led to a division of the IR image into 10 columns, each 68 pixels wide. From Figure 1-6, it is seen that Median Filter requires five (5) lines of delay, Gradient Operator requires eight (8), and Non Maximum Suppression requires seven (7) for a total of 20, since the algorithms operate sequentially. The total number of shifts is 1360 per column. At a clock frequency of 100 kHz, numerical degradation in the order of 20% will occur, which is probably too high. The MTF can be reduced in several ways.

The modulation transfer function is a function of the input signal frequency, the frequency of the shifts (clock frequency), the number of shifts, and the transfer efficiency. The more practical avenues of reduction are clock frequency and the number of shifts; we can double the number of operators to 20 each, and halve the clock frequency and number of shifts to 50 kHz and 680, respectively. This may produce an improvement to 10% degradation, but this number would have to be confirmed experimentally. Of course this

approach increases the total chip area which is still small and the external clocking circuitry. Operating at cryogenic temperatures will probably increase the transfer efficiency somewhat. Further, the input frequency can be band-limited to decrease the MTF. Using the 20 column segmentation, each 34 pixels wide, a total of 680 shifts are required to perform the Median Filter, Gradient Operator, and Non Maximum Suppression Algorithms at a 1 megapixel/sec. rate.

Moreover, surface channel CCD's are suitable for this task within the defined operating parameters, and the advantage of these devices is realizing the non destructive taps. These taps are necessary in extracting the appropriate pixels for the moving windows discussed in Section 1.3 and the second quarterly report.

The size required for achieving a memory 680 elements long is 1000 square mils if four phase clocking is employed. Hence for 20 columns we will require a silicon area 1000 mils long by 20 mils wide.

Operation of the memory at cryogenic temperatures will present no problems since its construction is similar to the other focal plane signal processing components. Considering the size of this memory chip we expect a yield of about 5%. The clocking circuits required for the memory module operation are not included in the area calculations.

3.4 General Observations

In this report we have considered the physical parameters of the focal plane signal processing circuit elements. In our opinion, no signal processing operation defined and discussed contains any inherent characteristics which will prevent fabrication. However, the number and size of the required IC

modules is considerable. Integration of all the aforementioned elements in a single large integrated circuit is a high risk effort. Development of each single IC block first should provide sufficient test vehicles and data needed to evaluate each signal processing component. This data should be obtained before any large scale integration of all the focal plane signal processing is undertaken. Such an approach will result in the most efficient method leading towards LSI focal plane signal processing.

4.0 FOCAL PLANE AREA

This section presents a preliminary estimate of the focal plane area occupied by the first portion of a cueing system. The estimate is preliminary in the sense that none of the clocking circuitry has been included in area estimates for the operators. The reason is that the methods by which the algorithms will handle the image's edges have not been specified. The estimate includes part of the left branch of the System Flow Chart of Figure 1-1, i.e., Median Filter, Gradient Operator, and the Serpentine Memory required for all three operators.

Assuming that the focal plane is divided into 20 columns, Table 4-1 shows the number of processors required for a system data rate of 1 mega-pixel/sec. It also shows the geometric area required for each processor and an estimate of the area as defined above. The area thus far is 1 inch \times 1/2 inch.

TABLE 4-1. PRELIMINARY ESTIMATE OF FOCAL PLANE AREA

<u>PROCESS</u>	<u>NUMBER REQUIRED</u>	<u>UNIT AREA</u>	<u>TOTAL AREA</u>
Serpentine Delay (34 pixels long)	400	1000 mils x 1 mil for 20	1000 mils x 20 mils
Median Filter	20	100 mils x 128 mils for 1	400 mils x 640 mils
Gradient Operator	20	8 mils x 10 mils for 1	32 mils x 50 mils
Non Maximum Suppression	20		
Frame Storage	2		
Line Storage (640 pixels long)	8		
Connected Components	4		
			1 in. x 1/2 in.